katacontainers

The speed of containers,

the security of VMs.

Kata Containers is an open-source container runtime, building lightweight virtual machines that seamlessly plug into the container's ecosystem.

Kata Containers is an alternative OCI compatible runtime that enhances the security of container workloads in a lightweight virtual machine.

PROJECT LAUNCH

December 2017

HYPER.SH

intel

HUAWEI

TODAY

99Cloud · Alibaba Cloud · ANT GROUP · arm · ATLASSIAN

Baidu 百度 · branch · 中国移动 China Mobile · CRAY · DELL Technologies · ebay

Google · HUAWEI · IBM · inspur 浪潮 · intel · INTERdynamix · Microsoft · NVIDIA

ORACLE · orange™ · Red Hat · SUSE · Tencent · VEXXHOST · ZTE

katacontainers

katacontainers.io

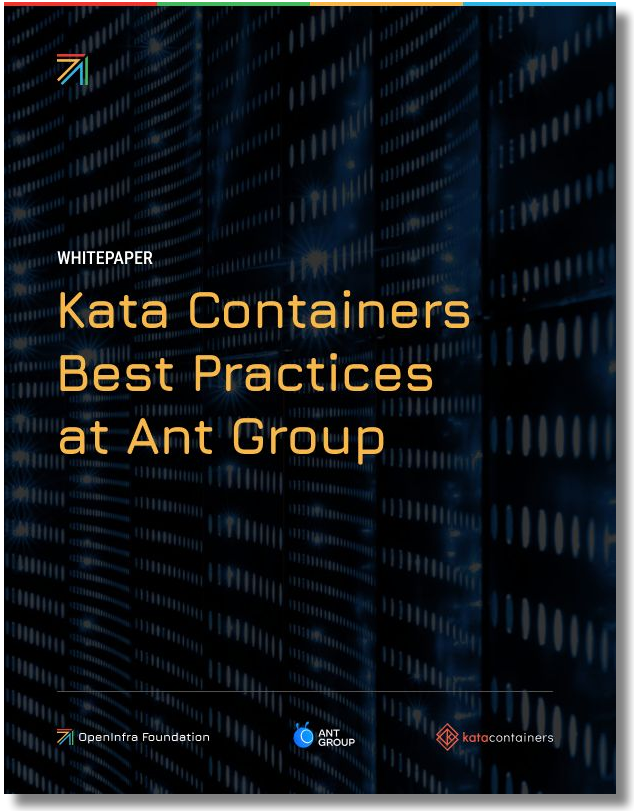

katacontainers.io

SUPERUSER

AWARDS

ANT GROUP

OVHcloud

Highlights

OCI-compatible runtime that enhances the security of container workloads in a lightweight virtual machines.

Works seamlessly with Kubernetes and Docker
and is a drop-in replacement for runc

Open Source
Open governance project under the Open Infrastructure Foundation umbrella

Multi Architecture
x86, ARM, IBM Power, IBM s/390x

Multi Hypervisor
QEMU,
Cloud Hypervisor,
Firecracker

Evolution

the speed of containers,
the security of VMs.

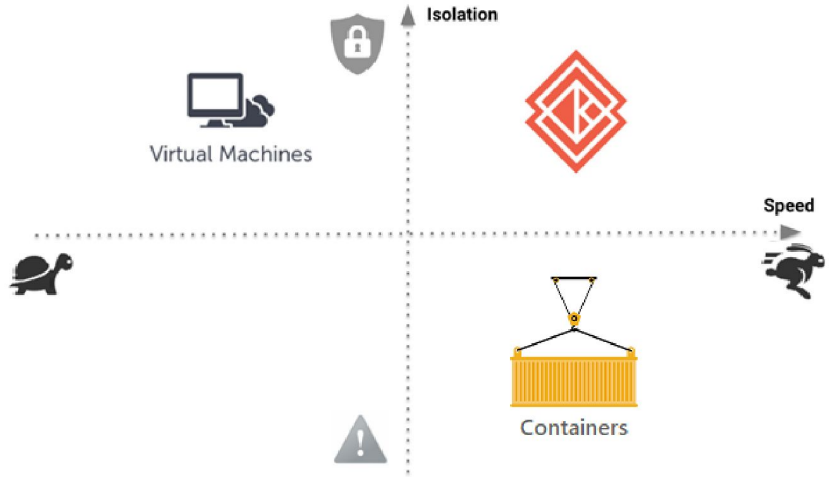

- 2015 – Intel launches Clear Containers open source project
- 2017 - Merger of two established projects under Open Infrastructure Foundation; Hyper.SH runV and Intel® Clear Containers.
- May 2018 – V1.0 released
  - Each container/pod isolated by a quick-to-boot lightweight VM.
  - OCI-compatible runtime - Looks just like a container in Kubernetes, Docker, or OpenStack.
- 2019 – Alibaba, Tencent, Baidu, Huawei, Stackpath put Kata Containers in production
- October 2020 – V2.0 released

Additional Security

Cgroups
Namespaces
Capability Filters
Seccomp filtering
Mandatory Access Control (MAC)

Standard Containers

+

Separate Guest Kernel
VMX non-root
Hardware control
CPU Access
Memory Access
Device Access

Virtual Machines

=

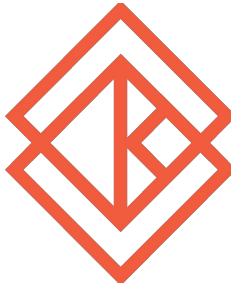

Kata Containers

Who is involved

Architecture Committee

Ant Group, Apple, Intel, Red Hat

Who's contributing or using?

Adobe, Alibaba, AMD, ARM, Atlassian, Baidu, Canonical, Google, eBay, Google, IBM, Inspur, Microsoft, NVIDIA, Oracle, Orange, Vexxhost, ZTE and other organizations

Who's interested?

Cloud service providers (IaaS, CaaS, etc.), network equipment makers (NFV workloads), banks and insurance companies, healthcare solution providers

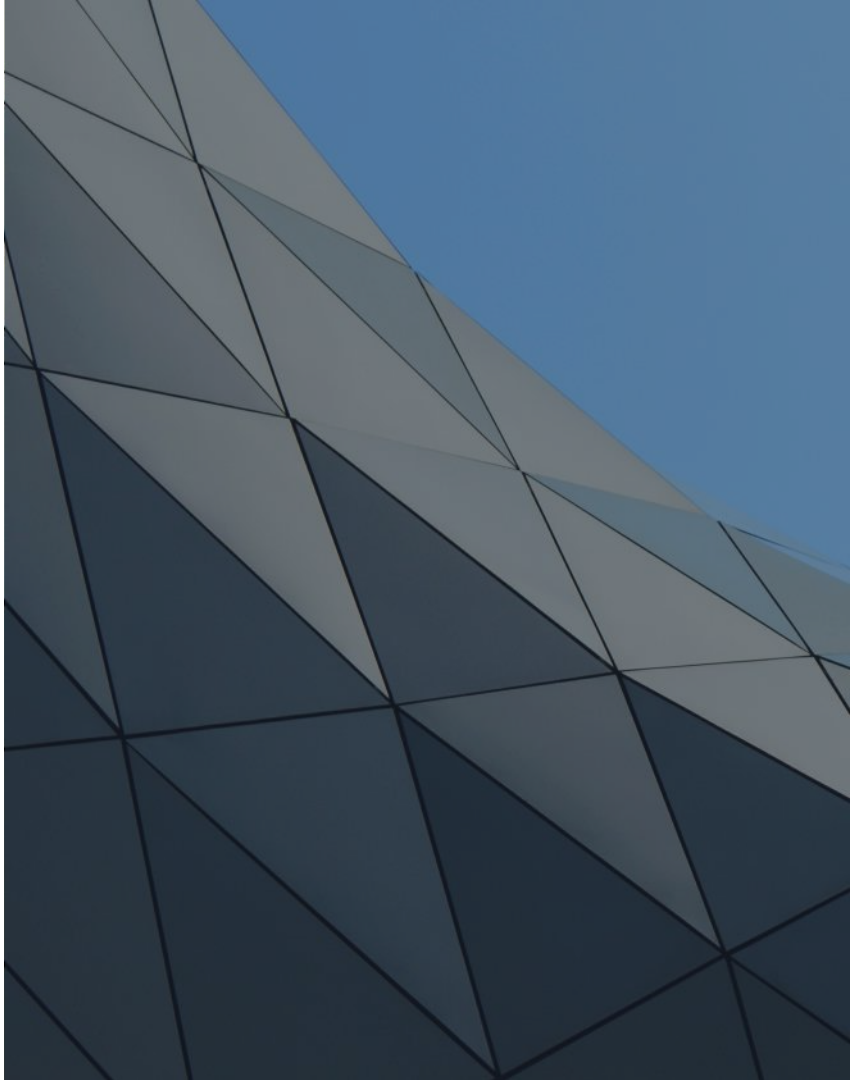

Healthy Growing Community

How does Kata Containers work?

How it works

Traditional Containers

Isolation by namespaces, cgroups with shared kernel

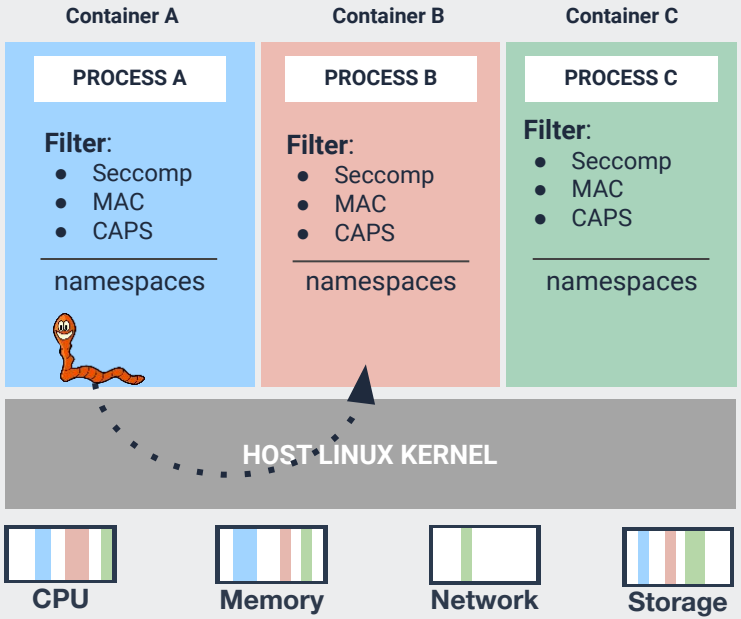

Kata Containers

Each container or pod is more isolated in its own lightweight VM

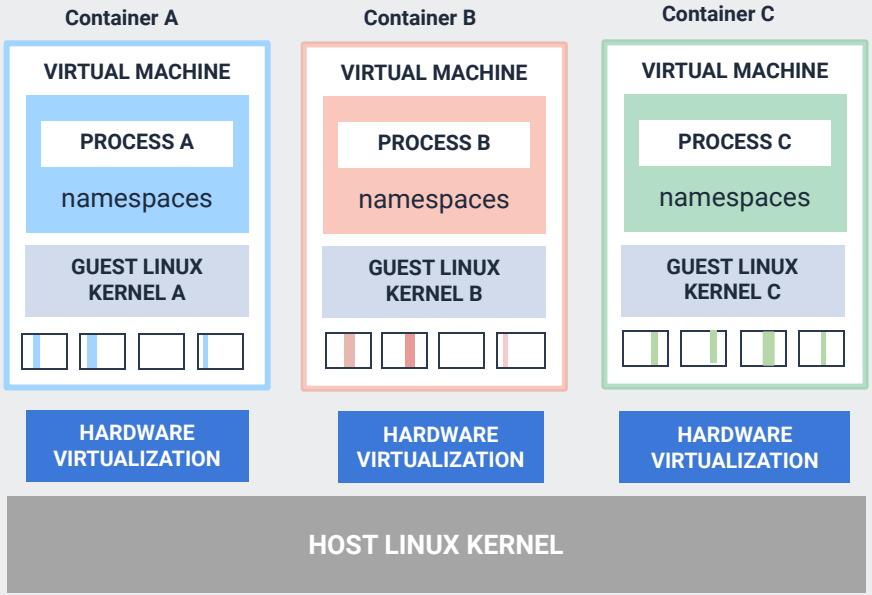

Seamless integration

Standard container initiation

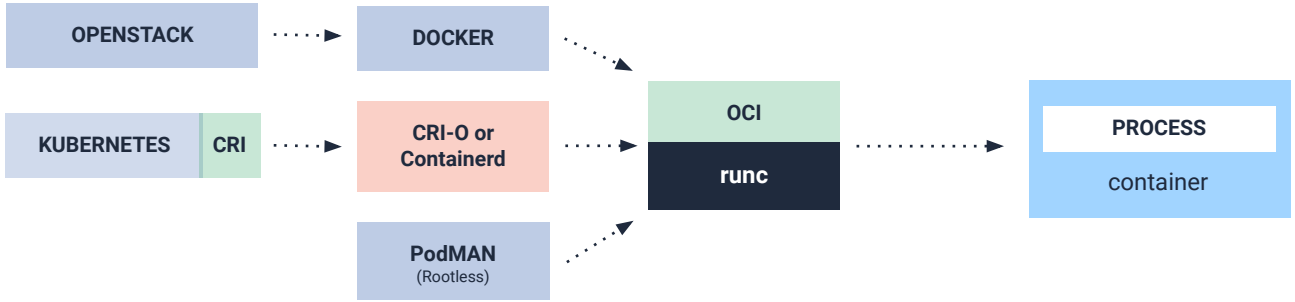

Kata Containers initiation

Looks and acts like a container using K8s, Docker, or OpenStack Zun

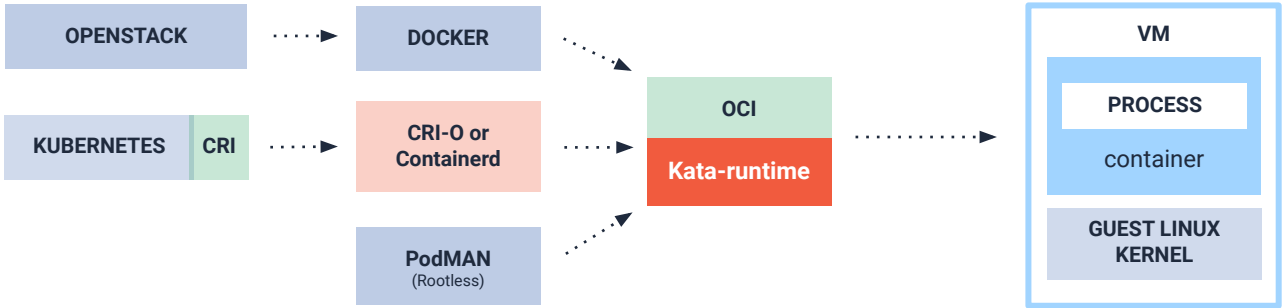

Multi-tenant Kubernetes

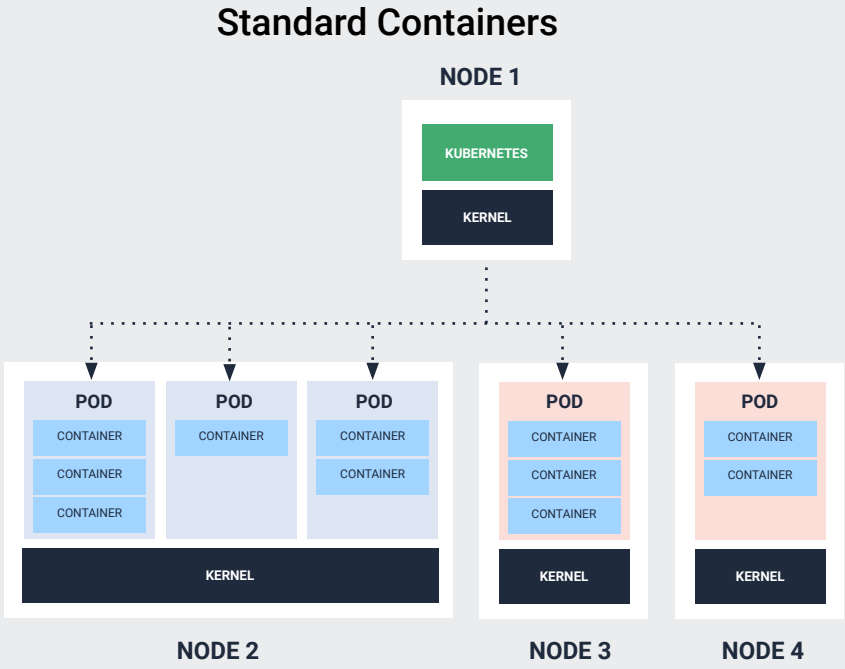

Isolate sensitive workloads by node

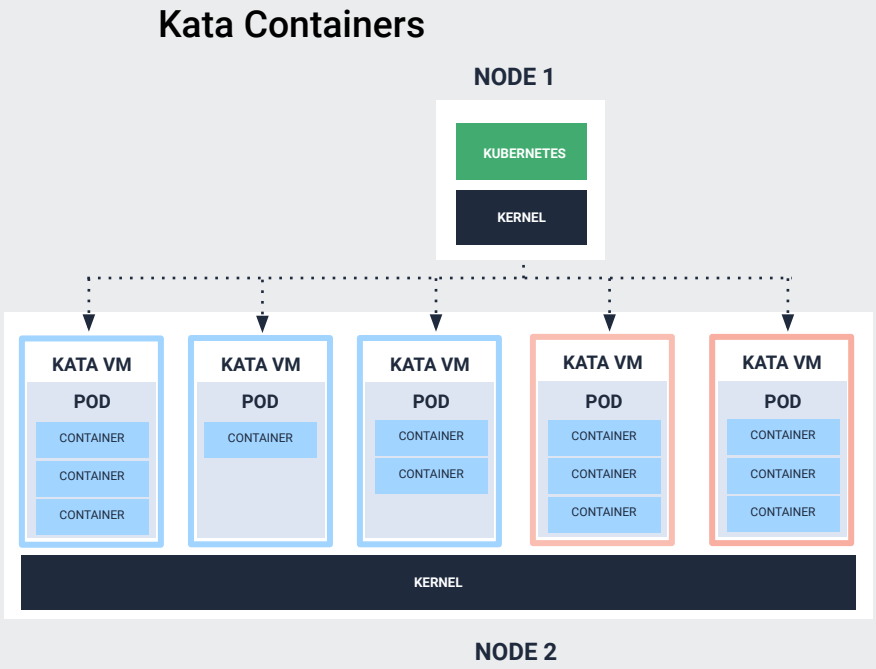

Isolate sensitive workloads **within** a node

Kata Containers “Sweet Spots”

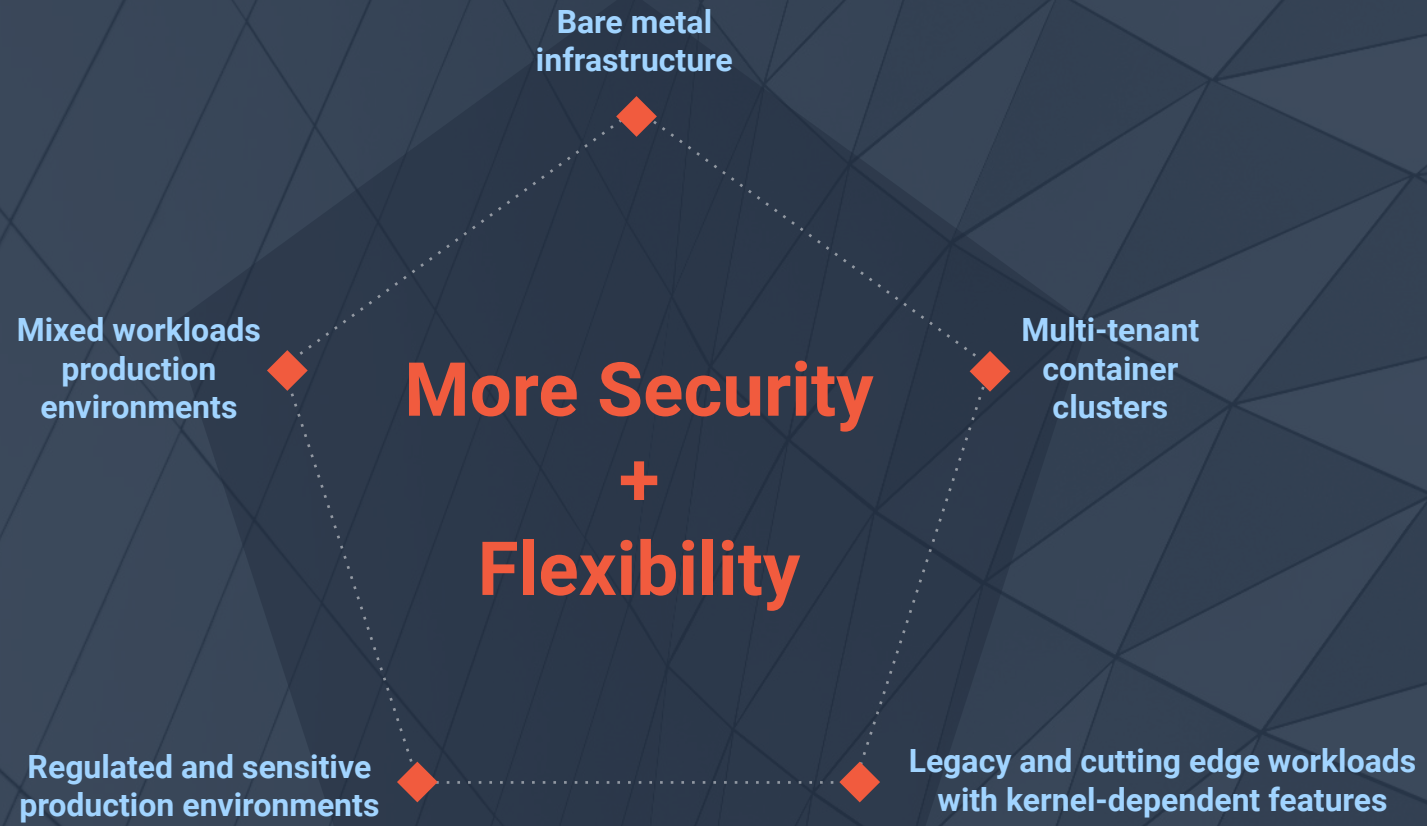

Features

Networking

SR-IOV (Lowest Latency)
Data Plane Development Kit (DPDK)/VPP (Fast Software implementation)
Network Plugins - CNI / CNM

Storage

SSD's / HDD
Virtio-fs (Shared filesystem) New
Virtio-SCSI/Virtio-blk (Block storage)

Intel® Optane™ Memory
DAX (Direct Access to Memory)

Memory

Kernel same page merging (Deduplicate Memory)
Virtio-mem (Experimental)

Hypervisor

QEMU
PCI device passthrough (Direct Device Assignment)
Hotplug of memory
DAX / NVDIMM (Direct Access to Memory)

Cloud Hypervisor
VMM for running modern cloud workloads
Feature parity with QEMU with a smaller attack surface

Firecracker
Fast and minimal

Kata Containers VM

Minimal kernel and rootfs (Customizable)
VM Templating (Fast Restore)

Kata Containers 2.0

Performance

- Transition to Rust Agent for higher container density and better memory overhead.
  11MB to 300K
- Transition from **gRPC** communication protocol to ***ttRPC*** for lower memory overhead.
- Virtio-fs is now the default shared file system type with better POSIX compliance.

Security

- Secure Enclave Support with Intel® SGX
- Reduced system privileges of Kata components
- Separate IO streams for better security isolation
- Cloud Hypervisor for smaller surface area of attack.

Stability

- New component called Kata-monitor
- Better observability and tracing for debug
- Live monitoring with Prometheus and Grafana.

Check https://github.com/kata-containers/kata-containers/releases/tag/2.0.0 for more information about Kata 2.0 and its download availability.

Where to run Kata Containers

Distro packages

Clear Linux
CentOS
Debian
Fedora
OpenSUSE
SUSE Linux Enterprise Server
Red Hat Enterprise Linux
Ubuntu

Cloud

Amazon Web Services (AWS)
Microsoft Azure
Google Compute Engine (GCE)
VEXXHOST OpenStack Cloud
Packet.IO

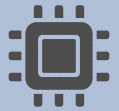

Hardware

Intel® architecture X86

AMD X86

ARM aarch64

IBM Z

IBM pSeries

The Kata Containers community welcomes contributions from anyone.

Go to katacontainers.io

Kata Containers Governance

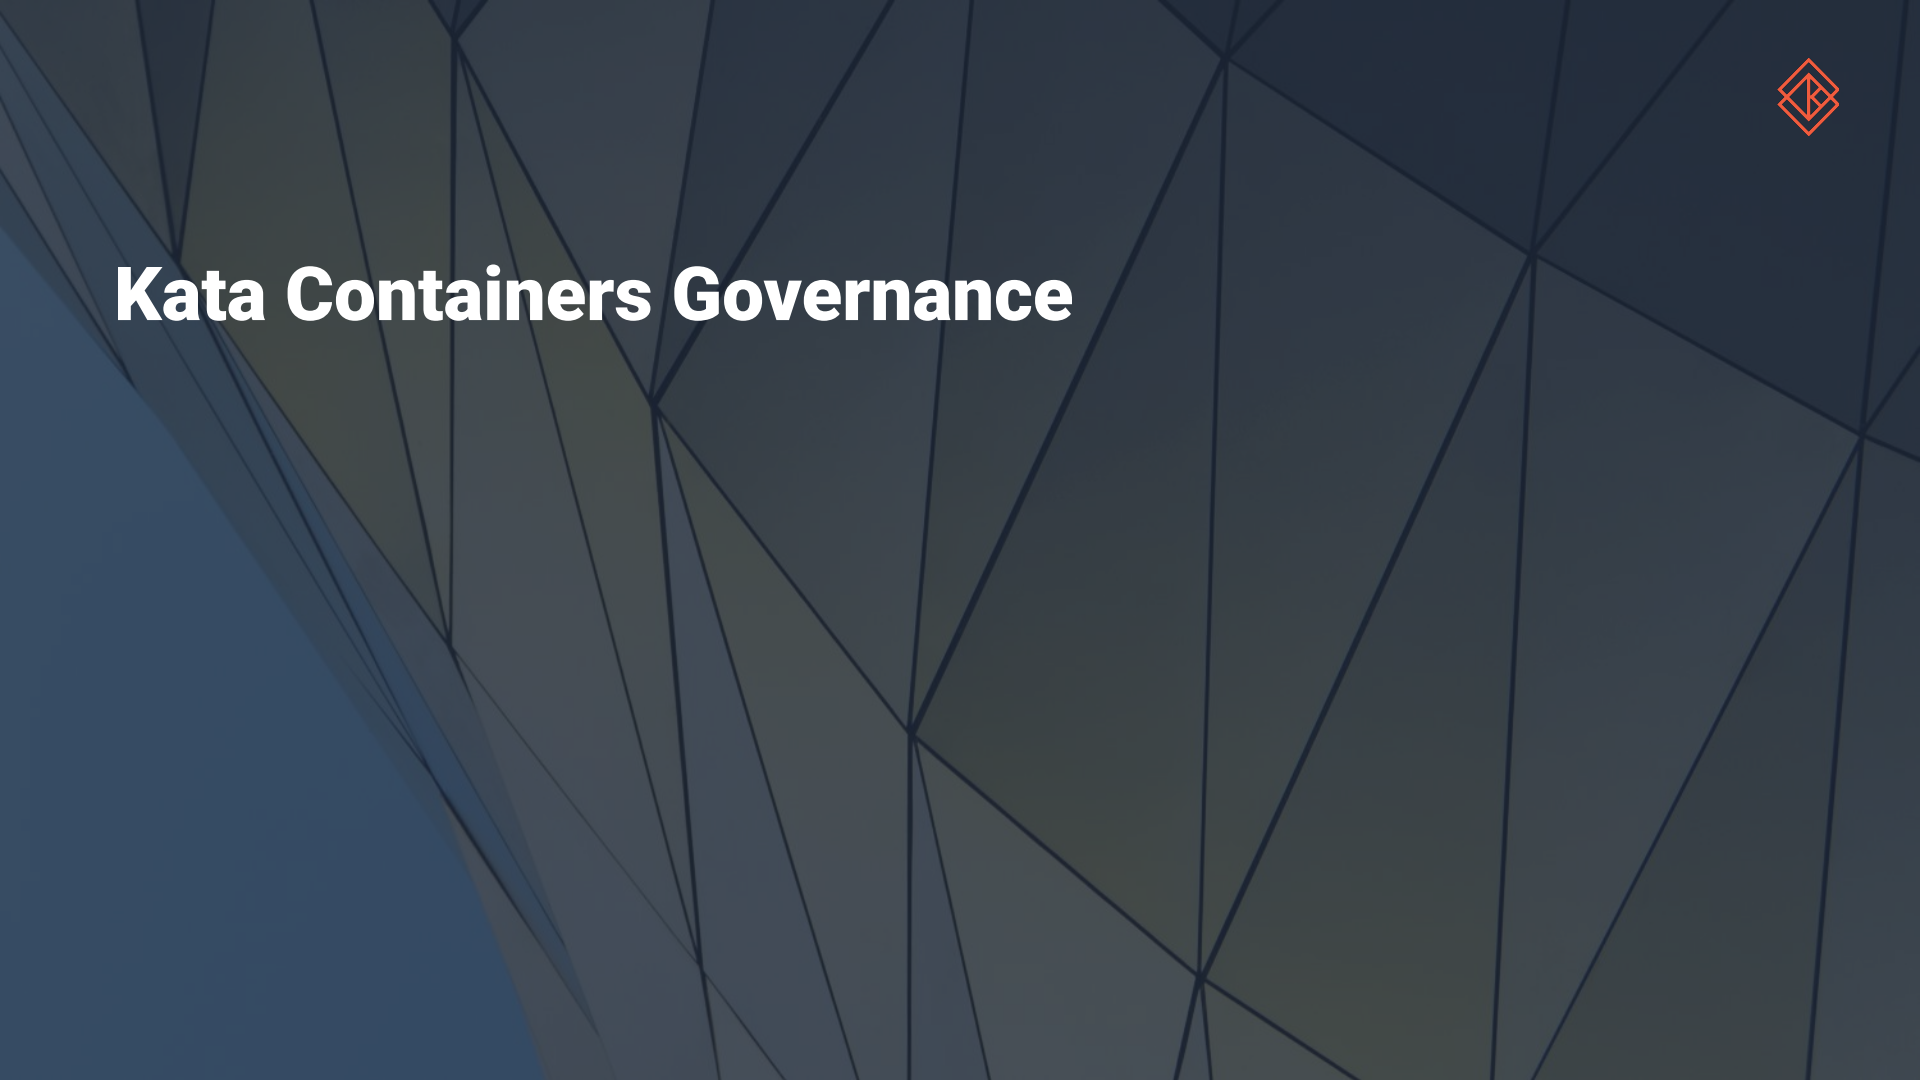

Governance

The Kata Containers project is governed according to the “four opens,”

- open source
- open design
- open development
- open community

Technical decisions will be made by technical contributors and a representative Architecture Committee. The community is committed to diversity, openness, encouraging new contributors and leaders to rise.

Governance

- **Contributors**
  - At least one github contribution for the past 12 months
- **Maintainers**
  - Active contributor, nominated by fellow maintainers
  - Can merge code
- **Architecture Committee**
  - Take high level architecture and roadmap decisions
  - 5 seats, elected by contributors

Governance

Architecture Committee

- The Architecture Committee is responsible for architectural decisions, including standardization, and making final decisions if Maintainers disagree.
- It will be comprised of 5 members, who are appointed by the Maintainers at launch but fully elected by Contributors within the first year.
- The Current Architecture committee members are Samuel Ortiz (Intel), Xu Wang (Ant Group), Eric Ernst (Apple), Archana Shinde (Intel), Fabiano Fidêncio (Red Hat)

Thank You!